NASA-CR-194334

RIACS

Research Institute for Advanced Computer Science
NASA Ames Research Center

USRA / RIACS

27 p

FINAL REPORT
ON
COOPERATIVE AGREEMENT NCC 2-387

JULY 1, 1992 - DECEMBER 31, 1992

Submitted to:
OFFICE OF UNIVERSITY AFFAIRS
NASA AMES RESEARCH CENTER

Submitted by:
Research Institute for Advanced Computer Science (RIACS)
An Institute of:
Universities Space Research Association (USRA)

RIACS Principal Investigator
JOSEPH OLIGER

NASA Technical Monitor
FRANK RON BAILEY

NASA Alternate Technical Monitor
MASAYUKI OMURA

Joseph Oliger

JOSEPH OLIGER, Director

N94-13932

(NASA-CR-194334) USRA/RIACS Final Report, 1 Jul. - 31 Dec. 1992 (Research Inst. for Advanced Computer Science) 27 p

Unclas

G3/61 0185487

USRA /RIACS

FINAL REPORT

ON

COOPERATIVE AGREEMENT NCC 2-387

JULY 1, 1992 - DECEMBER 31, 1992

Submitted to:
OFFICE OF UNIVERSITY AFFAIRS
NASA AMES RESEARCH CENTER

Submitted by:
Research Institute for Advanced Computer Science (RIACS)

An Institute of:
Universities Space Research Association (USRA)

R I A C S

JULY 1 - DECEMBER 31, 1992

C O N T E N T S

I.	INTRODUCTION	1
II.	RESEARCH IN PROGRESS	3
III.	TECHNICAL REPORTS	15
IV.	RIACS PUBLICATIONS	17
V.	SEMINARS AND COLLOQUIA	19
VI.	OTHER ACTIVITIES	20
VII.	RIACS CURRENT STAFF	23

I. INTRODUCTION

Research Institute for Advanced Computer Science

JOSEPH OLIGER
Director

The Research Institute for Advanced Computer Science (RIACS) was established by the Universities Space Research Association (USRA) at the NASA Ames Research Center (ARC) on June 6, 1983. RIACS is privately operated by USRA, a consortium of universities with research programs in the aerospace sciences, under a cooperative agreement with NASA.

The primary mission of RIACS is to provide research and expertise in computer science and scientific computing to support the scientific missions of NASA ARC. The research carried out at RIACS must change its emphasis from year to year in response to NASA ARC's changing needs and technological opportunities. A flexible scientific staff is provided through a university faculty visitor program, a post doctoral program, and a student visitor program. Not only does this provide appropriate expertise but it also introduces scientists outside of NASA to NASA problems. A small group of core RIACS staff provides continuity and interacts with an ARC technical monitor and scientific advisory group to determine the RIACS mission. RIACS activities are reviewed and monitored by a USRA advisory council and ARC technical monitor.

Research at RIACS is currently being done in the following areas:

- Parallel Computing
- Advanced Methods for Scientific Computing
- Learning Systems
- High Performance Networks and Technology
- Graphics, Visualization and Virtual Environments

In the past year, parallel compiler techniques and adaptive numerical methods for flows in complicated geometries were identified as important problems to investigate for ARC's involvement in the Computational Grand Challenges of the next decade. One of RIACS staff scientists, Dr. Robert Schreiber, has devoted most of his time to parallel compilers, and a senior visitor, Professor Marsha Berger, who works on adaptive grid methods, was invited to visit RIACS from September 1, 1991 through August 31, 1992.

During the past six months Professor Petter Bjorstad, of the University of Bergen, has begun a visit to RIACS (7/1/92 - present). He will work on domain decomposition techniques for PDEs for use on parallel computers. Dr. Pelle Olsson, from the University of Uppsala, joined RIACS as a post doctoral scientist in November 1993. He is working on the derivation of boundary conditions for the Navier-Stokes equations which are stable for nonsmooth boundaries and their efficient implementation on parallel computers.

We concluded a summer student visitors program during this six months. We had six visiting graduate students that worked on projects over the summer and presented seminars on their work at the conclusion of their visits.

RIACS technical reports are usually preprints of manuscripts that have been submitted to research journals or conference proceedings. A list of these reports for the period July 1, 1992 through December 31, 1992 is in the Reports and Abstracts section of this report.

[illegible] 2 [illegible]

II. RESEARCH IN PROGRESS

A. PARALLEL COMPUTING

HIGH PERFORMANCE FORTRAN FORUM
Robert S. Schreiber

Schreiber is an active participant in the High Performance Fortran Forum (HPFF), an academic, government, industry group that is designing a portable language based on Fortran 90 for programming distributed memory parallel machines. The Forum has approximately 40 participants. The group has met every six weeks throughout 1992 in order to complete its work and present a draft of the language at the Supercomputing '92 Conference. Schreiber has been working on the issues of language compatibility with Fortran 77, on the semantics of the data distribution directives, and on the library of new intrinsic functions. He has been chairman of the Intrinsics Committee.

HPFF is completing its work on schedule. A draft version was presented at Supercomputing '92 and is now available for public comment via *ftp*. A final meeting in March 1993 will be used to finish the final draft.

A number of hardware and software vendors have announced their intention to develop and market HPF compilers. It appears that HPF will be accepted by the parallel computer industry and become a *de facto* standard. If this happens, the Forum's work will have been a success.

RESEARCH IN HIGHLY PARALLEL MATRIX COMPUTATION
Robert S. Schreiber

Robert S. Schreiber, with Sid Chatterjee and Edward Rothberg, is continuing to investigate scalable methods for sparse Cholesky factorization on distributed memory machines. These methods must use block distribution (rather than row or column distribution) of the sparse matrices. Load balance, communication load, and scheduling are all important issues in achieving high performance on highly parallel machines.

With Elizabeth Jessup, they are also looking into scalable methods for the dense, symmetric eigenvalue problem. To date, there has been no satisfactory method implemented for the reduction of such a matrix to tridiagonal form with orthogonal similarity transformations. Such a method could be the basis for a standard highly parallel eigenvalue solver, since parallel methods for the tridiagonal eigenvalue problem are now available.

With Petter Bjorstad, a software message routing system for grid and toroidal machines has been developed in order to assess their usability for solving PDEs with unstructured meshes. This work will continue during the coming six-month period.

Also, in progress: *Efficient, Massively Parallel Eigenvalue Computation*, with Yan Huo (Electrical Engineering Department, Princeton University).

GENERATING LOCAL ADDRESSES AND COMMUNICATION SETS FOR DATA-PARALLEL PROGRAMS

Robert Schreiber
Siddhartha Chatterjee
John R. Gilbert (Xerox PARC)
Fred J. E. Long (UC Santa Cruz/RIACS)
Shang-Hua Teng (Xerox PARC/MIT)

The data-parallel language High Performance Fortran provides a two-level scheme for mapping arrays on a distributed-memory parallel computer and allows general block-cyclic distribution of arrays. The combination of these features substantially complicates the two runtime issues of generating local addresses and generating communication sets. A technique based on finite-state machines that efficiently solves both problems has been developed. The technique has been prototyped on the iPSC/860 and has been demonstrated to be extremely efficient.

This technique is the first to handle the full mapping scheme in High Performance Fortran in a unified and efficient manner. While solutions to various special cases have been known and implemented, this is the first general solution of the problem. This solution is being studied by the compiler groups at several major US vendors of parallel machines for possible inclusion in their systems.

AUTOMATIC ARRAY ALIGNMENT FOR DISTRIBUTED-MEMORY PARALLEL COMPUTERS

Robert Schreiber
Siddhartha Chatterjee
John R. Gilbert (Xerox PARC)
Shang-Hua Teng (Xerox PARC/MIT)

Data-parallel languages like Fortran 90 express parallelism in the form of operations on data aggregates such as arrays. Misalignment of the operands of an array operation can reduce program performance on a parallel machine by requiring nonlocal data accesses and communication. Determining array alignments to reduce communication is therefore a key issue in compiling such languages.

Work is continuing on a framework for the automatic determination of array alignments in data-parallel languages that supports operations such as array sectioning, reductions, spreads, and transpositions. A general technique, called compact dynamic programming (CDP), has been developed for such optimizations. Alignment functions are decomposed into three constituents: axis, stride, and offset. For each of these subproblems, techniques have been developed to solve the alignment problem for a basic block of code, possibly containing common subexpressions. Alignments are generated for all array objects in the code, both named program variables and intermediate results. Several of the algorithms have been prototyped to demonstrate the validity of the approach and confirm that they can be implemented efficiently.

The latest extension to this theory has been in the area of dynamic alignments depending on loop induction variables. By allowing the stride and offset alignments of arrays to be affine functions of the induction variables of the loop nest within which they occur, techniques have been developed to determine alignment parameters that minimize residual communication.

MASSIVELY PARALLEL COMPUTERS AND COMPUTATIONAL ELECTROMAGNETICS

Niel K. Madsen

Madsen has completed the parallelization of the preprocessor, PREDSI3D, for the time domain electromagnetic modeling code DSI3D. Initial performance tests have indicated that some parts of the preprocessor perform well on the Intel iPSC/860 computer, while other parts perform less well. A number of the subroutines in PREDSI3D are required to perform disk read and write operations and we have found that disk i/o can be a real bottleneck on this particular computer. As there are only ten i/o nodes, when more than ten nodes need to perform i/o linear speedup with the number of processors does not occur. Methods are currently being investigated to minimize the i/o problems.

The basic data allocation scheme used to distribute the work among the processors has been changed. Previously, the actual problem field component variables were distributed. This led to the need to solve very large partitioning problems. Now the problem grid cells are distributed among the processors. As there are as many as six field component variables associated with each cell, this strategy reduces the magnitude of the partitioning problem by a corresponding factor of six. This change of partitioning strategy required relatively minor changes to the preprocessor.

As replacements for the existing parallel machines are imminent, the DSI3D codes will be ported to these new machines - which will in principle provide a significantly greater 3D modeling capability. As in the past, it is expected to challenge these new machines with unstructured grid electromagnetic problems which are larger than have been solved heretofore.

B. ADVANCED METHODS FOR SCIENTIFIC COMPUTING

ITERATIVE METHODS FOR NON-HERMITIAN LINEAR SYSTEMS

Noël M. Nachtigal

Nachtigal, together with Roland Freund at AT&T Bell Labs, continued work in the area of iterative methods for non-Hermitian linear systems. In particular, they were interested in developing an alternate implementation of the quasi-minimal residual (QMR) method, based on a coupled two-term recurrence version of the underlying nonsymmetric Lanczos algorithm, rather than the classical three-term Lanczos recurrence. The motivation behind this work is that, in finite precision arithmetic, algorithms based on coupled two-term recurrences seem to have better numerical properties than the mathematically equivalent algorithms based on three-term recurrences.

They succeeded in developing a coupled version of the Lanczos process, complete with a look-ahead strategy that is necessary to avoid breakdowns. They have then applied the quasi-minimal residual approach, obtaining a new implementation of the QMR algorithm. Preliminary numerical experiments indicate that the new version does indeed have numerical properties that are superior to those of the original QMR algorithm.

Work is in progress on writing a software package that implements the new version of the QMR algorithm for the solution of non-Hermitian linear systems. When finished, it is planned to make this software available to the scientific community, as has been done in the past with the original QMR implementation. As an aside, Thinking Machines Corporation has ported the codes made available with the original QMR algorithm to the Connection Machine and has incorporated this version in their CM Scientific Software Library.

FINITE DIFFERENCE APPROXIMATIONS OF SYMMETRIC HYPERBOLIC AND PARABOLIC EQUATIONS IN SEVERAL SPACE DIMENSIONS

Pelle Olsson

The energy method has been applied to linear symmetric hyperbolic and parabolic systems in several space dimensions to derive stable gradient-free boundary conditions, which are suitable for implementation on a dataparallel computer. Boundary conditions leading to increased communication should be avoided whenever possible.

Existing stability results usually rely on the assumption of smooth domains. It is shown how to obtain stability results for linear symmetric hyperbolic and parabolic systems on nonsmooth domains. As an illustration, consider the linearized Euler and Navier-Stokes equations, including the energy equation. Nonlinear boundary conditions are obtained by interpreting the variables of the linearized problem as increments of those of the nonlinear problem. The assumptions on the domain are fairly mild; boundedness is assumed, but it is not necessary that the domain boundary be smooth everywhere. Hence, corners are allowed.

Since finite difference methods are employed, the domain of definition of the PDE is assumed to be diffeomorphic to an n-dimensional hypercube. Stable semi-discrete schemes of arbitrary order of accuracy are derived using one-sided difference stencils at the boundaries. Stability for semi-discrete hyperbolic and parabolic systems is inferred from the discrete version of the energy method using generalized norms. For certain norms it is possible to prove stability on domains with nonsmooth boundaries. The effects of the boundary conditions are taken into account.

It is demonstrated that the boundary conditions may be viewed as projection operators. This observation forms the basis for the implementation on parallel computers of SIMD architecture. On a serial computer, the time for evaluating the boundary conditions is negligible compared to time needed for updating the interior points. This may not be the case on a parallel computer. Assuming that the number of boundary operations is of the same order of magnitude at each grid point, the total execution time for enforcing the boundary conditions will be proportional to the number of boundary portions that must be sequentially treated. Clearly, a highly irregular boundary may reduce the inherent concurrency. It has been shown, however, that the boundary conditions for parabolic and hyperbolic PDEs may be concurrently evaluated, even if the boundary is nonsmooth. This can be done using Fortran 90 for the Euler and Navier-Stokes equations.

Olsson has also derived boundary-modified, semi-definite artificial viscosity operators of arbitrary order of accuracy. The viscosity operators are presented in a form that is particularly well-suited for the implementation on dataparallel computers.

DYNAMIC ADAPTATION AND PARALLEL IMPLEMENTATION OF 3-D UNSTRUCTURED GRIDS TO HELICOPTER AERODYNAMICS AND ACOUSTICS

Rupak Biswas

This project, in collaboration with Roger Strawn, Code RFA of U.S. AFDD, ATCOM, at NASA Ames Research Center, involves the development of efficient techniques for the fast refinement and coarsening of three-dimensional unstructured grids. These methods will be implemented on parallel computers and used to solve realistic problems in helicopter aerodynamics and acoustics.

The unsteady 3-D Euler equations are solved for helicopter flowfields using an unstructured grid. Error indicators are used to identify regions of the mesh that require additional resolution. Similarly, regions with low errors are targeted for coarsening. The objective is to optimize the distribution of mesh points so that the flowfield is accurately modeled with a minimum of computational resources. The mesh coarsening

and refinement algorithm is the key to the success of this procedure. The data structure for this algorithm is implemented in *C* and consists of a series of linked lists. These linked lists allow the mesh connectivity to be rapidly reconstructed when individual mesh points are added and/or deleted. It also allows for anisotropic refinement of the mesh.

Aerodynamics calculations performed on structured-grids have difficulties resolving localized flowfield features such as shocks, vortices, and acoustic waves. Unstructured-grid models can make use of localized mesh refinement to resolve these flow features. However, this mesh refinement is only effective if it can be performed efficiently in three dimensions. This new procedure for dynamic mesh adaption directly addresses this problem by using an innovative data structure that is well suited for large-scale computations. When coupled with a 3-D unstructured-grid Euler solver, the mesh adaption scheme will provide accurate solutions for complex aerodynamic flowfields.

The dynamic mesh adaption scheme is currently being tested for large problems on a Cray Y-MP computer. Particular attention is focused on CPU time and memory requirements. Future work will test the performance of the mesh adaption scheme and 3-D Euler solver on a massively parallel computer system.

RESEARCH IN DOMAIN DECOMPOSITION

Petter Bjorstad
Barry Smith (UCLA)

This work consists of two main lines of development. Petter Bjorstad and Barry Smith (UCLA), are writing up a more complete description and discussion of various algorithms that have been published over the years (this work takes the form of a book manuscript). A significant part of this work is to recast algorithms into a consistent form and describe them such that their implementation on parallel computers becomes feasible for non-experts in the field. This is important in order to get this powerful class of algorithms used in real application work. Along with this work is also actual software engineering and parallel implementations, in particular targeted towards problems from CFD calculations (this research is also joint with Eric Grosse and Bill Coughran of Bell Labs).

State of the art iterative methods for the solution of large scale discrete systems of equations arising from the discretization of partial differential equations are the so called Krylov Space Methods. These methods require the efficient application of a preconditioner as well as the discrete operator acting on a vector. Large scale computations in 3 dimensions often make use of highly unstructured grids. It is, therefore, of interest to investigate if the necessary communication in a parallel computer can be carried out without a large extra cost in this case. This research is joint with Robert Schreiber and considers the development of a communication compiler for 2-D, mesh interconnected parallel computers. In particular, data parallel (SIMD-style) machines like the MasPar, are considered.

RESEARCH IN HIGHLY PARALLEL COMPUTATION

Petter Bjorstad
Rob Schreiber
Erik Boman (Stanford)

Petter Bjorstad and Rob Schreiber, have developed a software message routing system for grid and toroidal machines, in order to assess their usability for solving PDEs with unstructured meshes. They will be continuing this work during the coming six-month period.

Bjorstad, with Erik Boman and Rob Schreiber, has also initiated the development of a high performance, modular implementation of the important BLAS-module DGEMM for parallel, distributed memory computers.

ADAPTIVE REFINEMENT OF COMPOSITE CURVILINEAR GRIDS

Steven Suhr

A software system is being designed and implemented to manage the adaptive refinement of composite curvilinear grids for the approximate solution of time-dependent partial differential equations. With the simplifying assumption that the spatial domain has fixed boundaries, an initial grid is constructed using a fixed set of overlapping grids, which collectively conform to the boundaries and cover the domain. Refinement grids, aligned with the original base grids, are added to maintain accuracy as the solution evolves. This approach organizes the grids into a geometrically nested tree of connected components, and it explores the use of curvilinear stairstep refinement grids. The adaptive grid system will initially be applied to model problems in two space dimensions, but it can be extended compatibly to three dimensions and to more realistic problems.

In implementing the adaptive grid system, a partial implementation of the new Vorpal language is being used, taking advantage of Vorpal's support for data structures, abstract data types, structured external files, and modular program structure. As the adaptive grid system evolves and grows, it should also be able to use the future support in Vorpal for concurrency and interactivity.

C. LEARNING SYSTEMS

MODEL-BASED LEARNING

Wray Buntine

This research involves developing methods and tools for doing supervised learning and regression. Supervised learning is the process of analyzing data to construct a "classifier", which is used for predicting one of the discrete variables in a record from the other variables.

In October 1992 a new applied project was initiated with Marshall Space Flight Center. This involves developing and applying data analysis methods to diagnosis Space Shuttle Main Engines via high resolution spectral readings. A pilot data analysis study on the Aviation Safety Reporting System completed in October 1992 yielded promising results, and a proposal is being prepared for more work.

The IND tree learning package was submitted to COSMIC in October 1992 as Version 2.0. This fixed numerous bugs in the early release and improved the interface. A new beta test version was released in December 1992 that included rule learning and decision graphs.

BAYESIAN SURFACE RECONSTRUCTION FROM MULTIPLE IMAGES

Peter Cheeseman
John Stutz (NASA Code FIA)
Robin Hanson (NASA/Sterling)
Bob Kanefsky (NASA/Sterling)

The Bayesian Learning group is involved in a number of projects that use Bayesian (statistical) methods to solve complex inference problems for noisy data. Such problems frequently arise in NASA applications. In 1992, a project was begun to take multiple images of the same area and use the information in these images to build a high resolution composite image that represents the programs best estimate of what the true surface must be. Initially we used multiple images of the surface of Mars (taken by the Viking Orbiter) to construct a high resolution composite image that shows features that are not visible in any of the contributing images. Currently, this technique is being extended to more complicated images, such as the Voyager

images of Ganymede, where the curvature of the planetary surface must be taken into account.

The development of a procedure for combining multiple image information required the solution to a number of difficult problems. It was particularly important to find very accurate registrations of the individual images to the composite, correct for nonuniform gain in the camera, learn and correct for the point-speed function within the camera, and other corrections related to noisy pixels. Algorithms were found for solving these problems and also for taking into account the neighbor correlations. The results of this work are reported in a NASA tech note, and have been submitted for presentation at a conference.

AUTOCLASS - AN AUTOMATIC CLASSIFICATION SYSTEM

Peter Cheeseman
John Stutz (NASA Code FIA)
Robin Hanson (NASA/Sterling)
Will Taylor (NASA/Sterling)

For several years, this group has been developing probabilistic techniques for the automatic classification of data. This is a useful technique for exploratory data analysis. This research has been implemented and tested in the AutoClass program. AutoClass accepts as input a set of independent data cases, each described by the same set of attributes. Then, using appropriate prior expectations and a likelihood model describing possible attribute relationships, AutoClass seeks the most probable classification given the data. It does not need to be told how many classes are present-it infers this from the data.

Recent AutoClass developments include models for full and partial correlation, hierarchical interclass relationships, and angular data. These developments, and the mathematics behind them, are described in a detailed NASA Tech Report and a conference paper. Copies of this experimental AutoClass version have been distributed to a number of researchers. A production version is currently distributed by NASA's COSMIC software center. A massively parallel version is now running on the NAS Connection Machine. AutoClass continues to provide a useful tool both for internal use, and externally.

AutoClass was applied to LandSat image data. It was found that AutoClass gives reasonable looking classifications and that some classes have obvious interpretations. A NASA Tech Report was prepared describing AutoClass and a classification of an intensively studied area in Kansas (First ISILCP Field Experiment), and collaborators familiar with the area are being sought. This LandSat work did not use spatial correlations among adjacent pixels and was unable to integrate information from successive images. These limitations led to the most recent project-surface reconstruction from multiple images.

SPARSE DISTRIBUTED MEMORY

Pentti Kanerva

This is long-term basic research into the engineering principles by which the brain and the rest of the nervous system organize and process information. Uncovering these principles serves dual purposes: Building of machines with skills, and designing of systems to fit human operators. The resulting technology will have wide use in aviation and in manned and unmanned space exploration.

Sparse distributed memory provides natural entry into the research. It was developed as a mathematical and engineering model of human long-term memory, and it is at once an associative memory, a random-access memory, a sequence memory, an artificial neural network, and a model of the cerebellum. The memory stores long (e.g., 1,000-bit) words that need not be precise, and it is also addressed by such words. An object or a scene or a moment of experience can be represented in the memory by a single word that contains a large amount of information about the object or the scene or the moment, and the word can serve equally as data and as a memory pointer. Memory pointers with high information content are at least partly

responsible for the synthetic and holistic powers of the brain. The idea is new to computer science and needs exploring. It will have a significant effect on computing once we learn to make full use of it.

A formal comparison of sparse distributed memory to the cerebellar model of Marr and to the Cerebellar Model Arithmetic Computer (CMAC) of Albus was completed and presented at a conference.

D. HIGH PERFORMANCE NETWORK & TECHNOLOGY

HIGH-SPEED NETWORKING
Marjory J. Johnson

The objective of this project is to support the High Performance Computing and Communications Program (HPCCP).

Several organizations, including major technology companies, research organizations, universities, and government laboratories, are working to establish a gigabit testbed within the Bay Area. M. Johnson has been collaborating with persons from the Numerical Aerodynamic Simulation Systems Division (NAS) and the NASA Science Internet Project Office (NSIPO) to represent NASA in this effort. As reported earlier, the unifying focus for the testbed is a teleseminar application, which will provide a tool to facilitate information exchange and collaboration. As the testbed matures, it is expected to promote active collaboration between industry, academia, and government, which is a major objective of HPCCP.

During the past six months, in addition to participating in general planning activities for the testbed, M. Johnson coordinated the writing of a proposal that was submitted to the Corporation for National Research Initiatives (CNRI) to obtain official national recognition for the testbed as a gigabit testbed. Comments from the reviewing committee regarding the testbed and the teleseminar concept were favorable; however, the committee was not convinced that gigabit bandwidth is required to support the teleseminar application. They have requested more information about specific testbed applications that will require gigabit bandwidth. Since recognition as a gigabit testbed would be likely to lend credibility to future testbed activities, testbed participants from NASA Ames are currently preparing a proposal describing plans for a networked-mass-storage project that is clearly a gigabit application for the testbed. We expect to submit this proposal to CNRI in January of 1993.

In a second HPCCP-related activity, M. Johnson is collaborating with the University of Arizona on a project in the Earth System Science (ESS) area. A three-year grant has recently been awarded by NASA to the University of Arizona under the HPCC/ESS Program to fund the activity. The objective of the project, entitled "Content-based Query and Browse of Earth Science Imagery Databases using High Performance Computers and Networks," is to develop algorithms for browsing large remote databases for specific patterns. Only images containing those patterns will be transferred to the user's home site, thus minimizing image-data transfer over the network. The pattern-recognition algorithms that are used for the database search will be implemented in neural networks running on parallel computers. The Principal Investigator of this project is located at the University of Arizona; M. Johnson is a Co-Investigator. Justin Paola, a graduate student at the University of Arizona, initiated work on this project while he was employed at RIACS this summer.

SPACE STATION NETWORKING ANALYSIS
Marjory J. Johnson

The feasibility of ground-based interactive command and control of Space Station Freedom scientific payloads is heavily dependent on performance of the end-to-end communications system. The objective of this project is to identify bottlenecks in the end-to-end system with respect to command and control of the Centrifuge Facility Project (CFP). The CFP was selected for this analysis because the project has relatively

high data-rate requirements and also because the project office is conveniently located at NASA Ames.

During the past six months M. Johnson has developed a model of the end-to-end Space Station communications system, using the simulation package Network II.5. The performance of the end-to-end system is dependent on such factors as network protocols, scheduling restrictions, limitations on uplink bandwidth, commanding procedure, space-to-ground communication outages, anticipated level of interaction between a scientist and his spaceborne experiment, and safety considerations. During the next six months, M. Johnson will refine the model and use it to determine the effects of the above factors on system performance.

The sponsor for this activity is Space Station Headquarters. The work is being conducted in collaboration with the Spacecraft Data Systems Research Branch of the Information Sciences Division, which is involved in analyzing broader data-management issues and conducting testbed activities in support of Space Station Freedom. Information about the CFP is being provided by project personnel.

RIACS NATIONAL SOFTWARE EXCHANGE (NSE) PROJECT
Mike Raugh

The RIACS NSE Project continues supporting the HPCC Software Exchange. The Software Exchange (SE) Project, led by Barry Jacobs at Goddard Space Flight Center is near the end of its second year. Participants representing the major HPCC Program agencies -- DARPA, DOC/NIST, DOE, and NSF, as well as NASA -- are cooperating to make information of interest and relevance to HPCC Program software developers available on-line through a shared SE infrastructure. RIACS is working with Jacobs to further develop and apply the GSFC Logical Library and Logical Books systems as prototypes for an on-line library and an experimental uniform interface, respectively. The Logical Library and its main catalogue permit easy lookup and access of major information resources on the Internet, such as RECON, Netlib, GAMS, WAIS, Cugdus, and CASRD, to name a few. Approximately forty-five such resources are available at this time, about a dozen of which are available in the "books" format.

The SE Project, and the RIACS supporting role, undergoes continual revision as strategies evolve rapidly for improving and populating the Library. The basic structure of the Library and Books systems have begun to stabilize, while the tools themselves are developing at a quickening pace. Beta versions of client/servers for both Books and the Library are nearing completion and multimedia capabilities have been incorporated and demonstrated. The assigned role for RIACS in this period has been to continue work on Books as a uniform access mechanism, in particular to develop the initial version of the Library main catalogue and publish it as a Book, to continue developing tools ingesting published materials into the Books format, and to develop tools for Library maintenance functions. Once these tools themselves have stabilized, they will be reported and made available to other HPCCP scientists and programmers who wish to publish Books.

The primary objective for RIACS in this half-year has been to work with the DAVID Publishing Group at GSFC to refine the process of publishing on-line Books, to seek additional resources that are relevant to workers in the HPCC, and to develop a first edition of the Library's catalogue of major resources. The term "Books" in this context refers to a particular way of representing information in both VT100 and, most impressively, X-Windows formats.

Such Books are proving to be a versatile and powerful means of representing widely differing categories of data in a consistent format, including software, textual information, images, recorded sound, and bibliographies. Books are being used even to describe hierarchies of information, leading from the more general levels of a taxonomy to more particular levels. A secondary objective for RIACS throughout the project has been to develop software for Books publishing.

Some of the major accomplishments of the SE Project as a whole have been: Beta versions of Books and Library client/servers, installation in the library of more than forty databases in their native modes, respectively, development of a Books cross-index for the joint holding Netlib at ORNL and the Guide to Available Mathematical Software (GAMS) at NIST. The National Software Exchange Working Group (NSEWG), sponsored by NASA and DARPA to permit critiquing of the SE by participating technologists and prospective users, held its initial organizing meeting at USRA headquarters in Washington, DC, in late summer and has scheduled its first working meeting for February 1993.

Principal RIACS contributions to the project in this period have been: publication of the Netlib Packages Book, publication of the first edition of the Library catalogue of major resources as a Book, and continuing development of an initial system of software for books publication. The Library and Books systems have been demonstrated at RIACS to ARC librarians, scientists, and managers. A talk entitled "Towards an HPCC Internet Library: A Walk through the HPCCP SE Project Prototype" was presented at the NAS/RIACS Forum on October 30.

We had planned to report steps required to publish Books in this period, but the continuing rapid evolution of our publication procedures has rendered it unwise to do so. We repeat that milestone for the next period.

The milestones planned for this period are: to continue publishing Books at the target rate of one per month (Jacobs at GSFC has agreed to continue helping to identify the sources for these books), and to report software tools and steps required to publish in the Books format.

E. GRAPHICS, VISUALIZATION & VIRTUAL ENVIRONMENTS

VISUALIZATION FOR PLANETARY EXPLORATION (VPE)

Lewis E. Hitchner

The objective of this project is to design and develop computer software systems and hardware interfacing for the Code FLM Virtual Planetary Exploration project. The VPE project supports "virtual" exploration of planetary terrains via the use of NASA satellite imagery databases, interactive computer graphics scene visualization and animation, and "virtual reality" human interface devices. The system is being developed for use by planetary geology scientific investigators and planetary mission planners. It is being developed with flexibility and extensibility to easily accommodate alternative terrain databases (e.g., terrestrial as well as planetary) at varying scales (global to human scale), future hardware and software platforms, and future display technologies. All development has been done in a networked graphics super-workstation environment using UNIX and X11 windows software plus locally written code.

The past six months' work has focussed on porting of the VPE software to a new software and hardware environment on our recently purchased Silicon Graphics SkyWriter graphics workstation. This dual pipeline graphics system provides an order of magnitude increase in rendering speed (animation frame update rate, polygons/sec.) over the older computer system that has been used since the project's beginning. The new system provides two parallel graphics hardware pipelines, four CPU's, and up to five independent, parallel video outputs. The VPE system software has been drastically revised to incorporate all these features. The vendor's visual simulation library software, SGI Performer, that is now being used, provides software interfaces to the high performance graphics and the multiprocessing and shared memory features of the hardware.

The current VPE virtual environment software architecture incorporates the following features:

1. DISTRIBUTED PROCESSING:

a) A master process runs on a host workstation running an X11 server. This process initiates the slave and provides a medium speed operator's user interface to the slave process via a socket connection.

b) A slave process runs on the high performance SGI SkyWriter that polls time critical user input devices (head tracker and 6 DOF joystick), responds to operator messages sent via the socket connection, and continuously draws animation frames. This process provides the rapid (10 to 30 Hz) video update for viewing in a stereoscopic, head-mounted display.

2. CONCURRENT RENDER PROCESSING:

The SkyWriter rendering operation uses 3 concurrent processes (application, cull, and draw) running on up to 4 CPU's. Each animation frame is pipelined through the three stages such that frame updates (the draw process) always occur as fast as possible with no waiting between updates. In dual channel stereo mode, the cull, draw, graphics hardware, and video output pipeline for each channel run in parallel.

The VPE software was extended to provide multi-resolution terrain modeling using several level-of-detail (LOD) modeling criteria. For each region of terrain several polygon mesh models at many different sampling resolutions are stored in a quad tree. At each update frame analysis of each subregion of the terrain surface determines an appropriate model resolution to use when drawing the terrain. The LOD analysis exploits the quad tree hierarchy and employs three criteria for selecting the appropriate model: range from the viewpoint, angular distance from the center of the field-of-view, and a user specified metric for application related regions-of-interest. The LOD analysis considers all three factors weighted according to user specifications. The analysis/feedback loop can be adjusted to respond to system load and user preferences. This technique provides both a high frame update rate and high level-of-detail for terrain regions important to the user (close to the viewpoint, near the center of field of view, and areas critical to the application) at the sacrifice of low detail in regions of low interest for visual perception and application requirements. The increased performance using this technique has greatly enhanced the VPE system's ability to provide compelling virtual environment simulations. This result of this work is being presented in a conference proceedings in February 1993.

This project is sponsored by the Human Interface Research Branch (Code FLM) of the Aerospace Human Factors Division (Code FL), NASA Ames. Dr. Michael W. McGreevy, Code FLM, is the NASA collaborator.

F. OTHER PROJECTS

GLOBAL CHANGE RESEARCH
Richard Johnson

The Earth Sciences Division (ESD) at NASA Ames Research Center has been exploring the research options that could be appropriate for support by the NASA Office of Commercial Programs (OCP) and various private sector industries. Support has been provided to ESD for the development of an overall strategic plan and a set of project options for discussions with OCP. This plan and project options were presented to and discussed with OCP program managers at NASA Headquarters via a video conference in July. Two follow-up presentations to OCP personnel were made in August at NASA Ames. R. Johnson discussed the opportunities for research and commercialization in the Biogeochemical Dynamics research and technology arena.

The status of on-going interactions of the Earth Sciences Division with several universities, the Electric Power Research Institute, and industries was presented. The interest of OCP personnel in the projects presented was high and support for one or more of the projects in FY93 appears to be likely.

Support for the development of programmatic and management plans for the MEDSAT project at NASA Ames and for the Remotely Piloted Aircraft Program have continued at a low level.

III. TECHNICAL REPORTS

AMR on the CM-2
MARSHA J. BERGER
JEFF SALTZMAN (Los Alamos National Laboratory)
TR 92.16 August 1992 (17 pages)

Submitted to Applied Numerical Mathematics

We describe the development of a structured adaptive mesh algorithm (AMR) for the CM-2. We develop a data layout scheme that preserves coarse grids. On 8K of a 32K machine we achieve performance slightly less than 1 CPU of the Cray Y-MP. We apply our algorithm to an inviscid compressible flow problem.

Optimal Evaluation of Array Expressions on Massively Parallel Machines
SIDDHARTHA CHATTERJEE
JOHN R. GILBERT (Xerox Palo Alto Research Center)
ROBERT S. SCHREIBER
SHANG-HUA TENG (Xerox Palo Alto Research Center)
TR 92.17 September 1992 (36 pages)

To appear in the Proceedings of the Second Workshop on Languages, Compilers and Runtime Environments for Distributed Memory Multiprocessors, Boulder, CO., September 30 - October 2, 1992.

We investigate the problem of evaluating Fortran 90-style array expressions on a massively parallel distributed-memory machine. On such machines, an elementwise operation can be performed in unit time for arrays whose corresponding elements are in the same processor. If the arrays are not aligned in this manner, the cost of aligning them is part of the cost of evaluating the expression. The choice of where to perform the operation then affects this cost. We present a dynamic programming technique to solve this problem efficiently for a wide variety of interconnection schemes, including multidimensional grids and rings, hypercubes and fat-trees. We also consider expressions containing operations that change the shape of the arrays, and show that our approach extends naturally to handle this case.

Automatic Array Alignment in Data-Parallel Programs
SIDDHARTHA CHATTERJEE
JOHN R. GILBERT (Xerox Palo Alto Research Center)
ROBERT S. SCHREIBER
SHANG-HUA TENG (Xerox Palo Alto Research Center)
TR 92.18 October 1992 (13 pages)

To appear in the Proceedings of the Twentieth Annual ACM SIGACT/SIGPLAN Symposium on Principles of Programming Languages, Charleston, SC, January 1993.

Dataparallel languages like Fortran 90 express parallelism in the form of operations on data aggregates such as arrays. Misalignment of the operands of an array operation can reduce program performance on a distributed-memory parallel machine by requiring nonlocal data accesses. Determining array alignments that reduce communication is therefore a key issue in compiling such languages.

We present a framework for the automatic determination of array alignments in dataparallel languages such as Fortran 90. Our language model handles array sectioning, reductions, spreads, transpositions, and asked operations. We decompose alignment functions into three constituents: axis, stride, and offset. For each of these subproblems, we show how to solve the alignment problem for a basic block of code, possi-

bly containing common subexpressions. Alignments are generated for all array objects in the code, both named program variables and intermediate results. The alignments obtained by our algorithms are more general than those provided by the "owner-computes" rule.

Finally, we present some ideas for dealing with control flow, replication, and dynamic alignments that depend on loop induction variables.

Implementation Details of the Coupled QMR Algorithm

ROLAND W. FREUND
NOËL M. NACHTIGAL

TR 92.19 October 1992 (17 pages)

The original quasi-minimal residual method (QMR) relies on the three-term look-ahead Lanczos process to generate basis vectors for the underlying Krylov subspaces. However, empirical observations indicate that, in finite precision arithmetic, three-term vector recurrences are less robust than mathematically equivalent coupled two-term recurrences. Therefore, we recently proposed a new implementation of the QMR method based on a coupled two-term look-ahead Lanczos procedure. In the paper, we describe implementation details of this coupled QMR algorithm, and we present results of numerical experiments.

IV. PUBLICATIONS

RUPAK BISWAS

Applications of an Unstructured Adaptive-Grid Euler Solver to Helicopter Aerodynamics and Acoustics, R. C. Strawn, R. Biswas, and M. Garceau, Compendium of Abstracts of the NASA Computational Aerosciences (CAS) Conference, NASA Ames Research Center, Moffett Field, CA, August 18-20, 1992, pp. 66-67.

A New Procedure for Dynamic Adaption of Three-Dimensional Unstructured Grids, R. Biswas and R. C. Strawn, to appear in the Proceedings of the 31st AIAA Aerospace Sciences Meeting, Reno, NV, January 11-14, 1993.

PETTER BJORSTAD

Two Different Data-Parallel Implementations of the BLAS (with Tor Sorevik, University of Bergen). To appear in NATO proceedings from a meeting on parallel computing, Cetraro, Italy, June.

Dataparallel BLAS As a Basis for LAPACK on Massively Parallel Computers (with Tor Sorevik, University of Bergen). In NATO Advanced Study Institute, published by Kleuver 1993.

Large Scale Structural Analysis on Massively Parallel Computers (with Jeremy Cook, University of Bergen). In NATO Advanced Study Institute, published by Kleuver 1993.

SIDDHARTHA CHATTERJEE

Optimal Evaluation of Array Expressions on Massively Parallel Machines, by Siddhartha Chatterjee, John R. Gilbert (Xerox PARC), Robert Schreiber, and Shang-Hua Teng (Xerox PARC). Appears in the Proceedings of the Second Workshop on Languages, Compilers, and Runtime Environments for Distributed Memory Multiprocessors, Boulder, CO, September 1992. Also available as RIACS TR 92.17.

Automatic Array Alignment in Data-Parallel Programs, by Siddhartha Chatterjee, John R. Gilbert (Xerox PARC), Robert Schreiber, and Shang-Hua Teng (Xerox PARC). To appear in the Proceedings of the Twentieth Annual SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Charleston, SC, January 1993. Also available as RIACS TR 92.18.

Compiling Nested Data-Parallel Programs for Shared-Memory Multiprocessors, by Siddhartha Chatterjee. Accepted for publication in ACM Transactions on Programming Languages and Systems.

Generating Local Addresses and Communication Sets for Data-Parallel Programs, by Siddhartha Chatterjee, John R. Gilbert (Xerox PARC), Fred J. E. Long (UC Santa Cruz), Robert Schreiber, and Shang-Hua Teng (Xerox PARC). To appear in the Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, San Diego, CA.

Implementation of a Portable Nested Data-Parallel Language, by Guy E. Blelloch (CMU), Siddhartha Chatterjee, Jonathan Hardwick (CMU), Jay Sipelstein (CMU), and Marco Zagha (CMU). To appear in the Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, San Diego, CA.

LEWIS E. HITCHNER

Virtual Planetary Exploration: A Very Large Virtual Environment, ACM SIGGRAPH '92, course notes for tutorial on "Implementing Immersive Virtual Environments," Steve Brsyon, course chair, Chicago, IL; August 1992.

The NASA Ames Virtual Planetary Exploration Testbed, Proceedings of IEEE Wescon '92, Anaheim, CA,

November 1992.

PENTTI KANERVA

Associative-memory Models of the Cerebellum, In I. Alexander and J. Taylor (eds.), Artificial Neural Networks 2(1):23-34 (Proceedings ICANN'92, Brighton, UK). Amsterdam: Elsevier, 1992.

NOËL NACHTIGAL

An Implementation of the QMR Method Based on Coupled Two-Term Recurrences - R.W. Freund (AT&T Bell Labs) and N.M. Nachtigal. "Linear Algebra for Large Scale and Real-Time Applications," Proceedings of the NATO Advanced Study Institute, Leuven, Belgium, August 3-14, 1992.

PELLE OLSSON

High Order Finite Difference Methods on Non-Smooth Domains - to appear in Computer Methods in Applied Mechanics and Engineering.

Finite Difference Approximations of Symmetric Hyperbolic and Parabolic Equations in Several Space Dimensions - submitted for publication, SIAM Journal of Scientific and Statistical Computing.

V. RIACS SEMINARS AND COLLOQUIA

Justin D. Paola - RIACS/NASA Ames Research Center - University of Arizona — July 29
The Application of Neural Networks in the Classification of Multi-Spectral Satellite Imagery

Shang-Hua Teng - XEROX Palo Alto Research Center — August 11
A Geometric Approach to Mesh Partitioning

Marsha Berger - RIACS/NASA Ames Research Center - New York University — August 25
Issues for Adaptive Mesh Refinement: Parallelism and Geometry

Petter Bjorstad - RIACS/NASA Ames Research Center - University of Bergen — Sept. 8
Benchmarking of Scientific Computers, SLALOM - A Case Study

James M. Strzelec - RIACS/NASA Ames - Stanford University — Sept. 9
Iterative Methods: An Informal Comparison

Bruce Hendrickson - Sandia National Laboratories — October 26
Spectral Octasection for Mapping Parallel Computations

Mike Raugh - RIACS/NASA Ames Research Center — October 30
Towards an HPCC Internet Library -- An Informal Discussion of the HPCC Software Exchange Prototype

VI. OTHER ACTIVITIES

RUPAK BISWAS

Attended NASA Computational Aerosciences (CAS) Conference, NASA Ames Research Center, Moffett Field, CA, August 18-20, 1992.

PETTER BJORSTAD

Presented *Dataparallel BLAS As A Basis for LAPACK On Parallel Computers,* an invited talk, NATO Advanced Study Institute, Leuven, Belgium, August 3-14, 1992.

Presented *Large Scale Structural Analysis and Parallel Computing*, an invited talk, NATO Advanced Study Institute, Leuven, Belgium, August 3-14, 1992.

Presented *Domain Decomposition and Large-Scale Computing*, 6th. annual meeting of NORTIM, Trondheim, Norway, August 20, 1992.

Presented *Domain Decomposition Algorithms and Parallel Computing,* UNI-C, (Danish Supercomputer Center) Copenhagen, Denmark, August 25, 1992.

Presented *Benchmarking of Scientific Computers, SLALOM - A Case Study,* RIACS/NASA Ames Research Center, Moffett Field, CA, September 8, 1992.

Served on an international committee to evaluate the quality of research in parallel computing in Sweden, January 22-26, 1993.

WRAY BUNTINE

Presented *Probabilities and Computing,* Carnegie Mellon University, Department of Computer Science, November 16, 1992.

Presented *Probabilities and Computing*, AT&T Murray Hill Laboratories, November 30, 1992.

Presented *Tree Classification Software*, Technology 2002, Baltimore, December 1, 1992.

Presented *Learning from Data*, Tudor Investment Corp., Manhattan, December 4, 1992

SIDDHARTHA CHATTERJEE

Presented *Compilers for Portable Data-Parallel Programming* in the Distinguished Lecture Series, MasPar Computer Corporation, Sunnyvale, CA, March 5, 1992.

Attended the Gordon Research Conference on Software Tools and Libraries for Concurrent Supercomputing, Plymouth, NH, July 27 - 31, 1992.

Was the external member of Ph.D. thesis committee for Deb Banerjee, Department of Mathematics and Computer Science, Dartmouth College, Hanover, NH. Attended his defense on July 31, 1992.

Presented paper, *Optimal Evaluation of Array Expressions on Massively Parallel Machines*, at the Second Workshop on "Languages, Compilers and Run-Time Environments for Distributed-Memory Multiprocessors", Boulder, CO, September 30 - October 2, 1992.

Supervised Fred Long, a RIACS summer student from UC Santa Cruz, who implemented a fast local address generation algorithm for mapped arrays in High Performance Fortran developed by Chatterjee, Gilbert, Schreiber, and Teng.

LEWIS E. HITCHNER
Co-taught a tutorial on *Implementing Immersive Virtual Environments* at the annual ACM SIGGRAPH conference, Chicago, IL, August, 1992.

Visited the Alpha_1 Graphics Research Group, Computer Science Dept., Univ. of Utah, and presented a seminar on the NASA Ames virtual environment research, Salt Lake City, Utah, September 1992.

MARJORY J. JOHNSON
Member of the program committee and session chair for the 4th IFIP Conference on High Performance Networking, Liege, Belgium, December 1992.

PENTTI KANERVA
Plenary talk on *Associative-Memory Models of the Cerebellum* at the Second International Conference on Artificial Neural Networks (ICANN'92) in Brighton, UK, September 3-7.

NOËL M. NACHTIGAL
Presented *An Implementation of the QMR Method Based on Coupled Two-Term Recurrences*, for a conference at the NATO Advanced Study Institute, Leuven, Belgium, August 3-14, 1992.

Presented *An Implementation of the QMR Method Based on Coupled Two-Term Recurrences*, at the SIAM 40th Anniversary Meeting, Los Angeles, CA, July 20-24, 1992.

Presented *Iterative Methods in Linear Algebra*, at the Naval Postgraduate School, Monterey, CA, November 19, 1992.

Presented *An Implementation of the QMR Method Based on Coupled Two-Term Recurrences*, at the University of Karlsruhe, Germany, September 11, 1992.

Presented a seminar, *Iterative Methods for Non-Hermitian Linear Systems*, at the IBM Scientific Center, Heidelberg, Germany, September 10, 1992.

Participated in the five-day course on Large Scale Scientific Computation at the University of Bielefeld, Bielefeld, Germany, August 31 - September 4, 1992.

JOSEPH OLIGER
Elected to the Society of Industrial and Applied Mathematics Board of Trustees, 1993-1996.

Chairman, National Center for Atmospheric Research Scientific Computing Division Advisory Committee.

Member, Organizing Committee for IMA Summer Program on *Modeling, Mesh Generation and Adaptive Numerical Methods for Partial Differential Equations*, 5-23 July 1993.

ROBERT S. SCHREIBER
Presented *High Performance Fortran Intrinsics* at Workshop on High Performance Fortran, Supercomputing '92, November 17, 1992.

Presented *Are Sparse Matrices Poisonous for Highly Parallel Machines* for UNI-C, (Danish Supercomputer Center) Copenhagen, Denmark, July 1992.

Presented *Language and Compiler Grand Challenges*, an invited talk; Gordon Research Conference on Software Tools and Libraries for Concurrent Supercomputing, Plymouth, NH, July 27-31, 1992.

Presented *Automatic Array Alignment in Data-Parallel Languages;* SIAM National Meeting, Los Angeles, CA, July 1992.

Organized meeting, "Software for Highly Parallel Computing," SIAM National Meeting, Los Angeles, CA, July 1992.

Co-organized meeting with Anoop Gupta (Stanford University), "New Directions in Parallel Computer Architecture," SIAM National Meeting, Los Angeles, CA, July 1992.

With Peter Bjorstad, Schreiber has developed a software message routing system for grid and toroidal machines, in order to assess their usability for solving PDEs with unstructured meshes. This work will be continued during the coming six-month period.

Also in progress: The Suitability of RISC Microprocessors for Scientific Computing with David Bailey (NASA Ames Research Center).

VII. RIACS STAFF

ADMINISTRATIVE STAFF

Joseph Oliger, Director - Ph.D., Computer Science, University of Uppsala, Sweden, 1973. Numerical Methods for Partial Differential Equations (03/25/91 - present).

Frances B. Abel, Office and Financial Manager (5/5/88 - present).

Deanna M. Gearhart, Administrative Assistant II (5/9/88 - present).

Anne Kohutanycz, Administrative and Systems Assistant (1/21/85 - 8/20/92).

Evangeline Tanner, Administrative Assistant I (4/5/90 - present).

TECHNICAL SUPPORT STAFF

Bradley Christofferson, A.A. - Computer Information Systems, DeAnza College 1987. Computer Systems Manager (10/23/89 - 10/30/92)

Jessica Casillas, B.A. - Mathematics, San Francisco State University 1986. Computer Systems Manager (11/25/92 - present).

RIACS SCIENCE COUNCIL

Dr. David Cummings (Interim Convener), Executive Director, Universities Space Research Association, Columbia, MD.

Dr. Dennis B. Gannon, (Covener), Director, Center for Innovative Computer Applications (CICA), Indiana University, Bloomington, IN.

Dr. Joseph Flaherty, Chair, Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY.

Dr. Joseph Oliger (Ex-Officio), Director, Research Institute for Advanced Computer Science, NASA Ames Research Center, Moffett Field, CA.

Dr. James W. Demmel, Computer Science Division, University of California, Berkeley, CA.

Dr. David Gottlieb, Division of Applied Mathematics, Brown University, Providence, RI.

Dr. Kenneth W. Neves, Boeing Company, Seattle, WA.

Dr. Thomas H. Pulliam, NASA Ames Research Center, Moffett Field, CA.

Dr. Daniel A. Reed, Department of Computer Science, University of Illinois, Urbana, IL.

Dr. Robert B. Schnabel, Department of Computer Science, University of Colorado, Boulder, CO.

Dr. Marc Snir, IBM Thomas J. Watson Research Center, Yorktown Heights, NY.

SCIENTIFIC STAFF

Wray Buntine, Ph.D., Computer Science, University of Technology, Sydney, Australia, 1992. Mathematical and probabilistic modeling of problems in intelligent systems (1/2/90 - present).

Peter Cheeseman, Ph.D., Artificial Intelligence, Monash University, Australia 1979. Artificial intelligence and automatic control, induction of models under uncertainty, Bayesian inference, expert systems and robotics (5/6/85 - present).

Dave Gehrt, JD Law, University of Washington, 1972. UNIX system administration, security, and network based tools (1/84 - 7/85, 2/88 - present).

Lewis E. Hitchner, Ph.D., Computer Science, University of Utah, 1984. Virtual environments, image synthesis/analysis, human/computer interaction, digital image processing and data compression, interactive graphics software systems (10/2/89 - present).

Marjory J. Johnson, Ph.D., Mathematics, University of Iowa, 1970. High-performance networking for both space and ground applications (1/9/84 -present).

Richard G. Johnson, Ph.D. - Physics, Indiana University, 1956. Global environmental problems and issues (1/4/88 - present).

Pentti Kanerva, Ph.D. - Philosophy, Stanford, 1984. Sparse distributed memory, autonomous learning systems (10/15/85 - present).

Eugene Levin, Ph.D. - Mathematics, UCLA, 1955. Computational Chemistry (6/6/83 - 9/25/92).

Michael R. Raugh, Ph.D. - Mathematics, Stanford University, 1977. Mathematics and computers for modeling physical and biological systems (1/28/85 - present).

Robert Schreiber, Ph.D. - Computer Science, Yale University, 1977. Parallel numerical algorithms and parallel computer architectures (8/29/88 - present).

VISITING SCIENTISTS

Marsha Berger, Ph.D. - New York University. Computational fluid dynamics; parallel computing (9/1/91-8/31/92).

Petter Bjorstad, Ph.D - Professor of Computer Science, University of Bergen, Norway. Parallel Computing, Domain Decomposition Algorithms for PDE's (7/1/92 - present).

POST-DOCTORAL SCIENTISTS

Rupak Biswas, Ph.D. - Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY. Large-scale scientific computation using parallel and adaptive methods (9/16/91 - present).

Siddhartha Chatterjee, Ph.D. - Computer Science, Carnegie Mellon. Compilation for distributed-memory parallel machines; parallel algorithms and applications (11/1/91 - present).

Noël M. Nachtigal, Ph.D. - Massachusetts Institute of Technology. Iterative methods for large, sparse, non-Hermitian linear systems (9/2 /91 - present).

Pelle Olsson, Ph.D. - Uppsala University, Sweden. Initial-boundary value problems for hyperbolic and

parabolic PDEs, numerical methods for PDEs on parallel computers (11/2/92 - present).

STUDENTS

Fred Long - Computer Science, UC San Diego, (1990). Construction, implementation and evaluation of algorithms for efficient address generation for FORTRAN 90 programs mapped onto distributed-memory parallel computers (6/22/92 - 8/28/92).

Justin Paola - Electrical Engineering & Computer Science, UC Berkeley, (1990). Image processing/ remote sensing (5/26/92 - 8/7/92).

Brian Rogoff - Mathematics, M.I.T. (1986). Scientific visualization (7/1/92 - 9/31/92).

James M. Strzelec - Applied Mathematics, SUNY, Stony Brook, (1992). Iterative methods (6/22/92 - 9/9/92).

Steven Suhr - Computer Science, Stanford University, (1987). Algorithms for scientific computing; programming languages (7/1/92 - same).

CONSULTANTS

Marsha Berger - New York University. Computational fluid dynamics; parallel computing (9/1/92 - present).

Tony F. Chan - Professor of Mathematics, University of California, Los Angeles. Efficient algorithms in large-scale scientific computing, parallel algorithms and computational fluid dynamics (10/01/86 - present).

Michael Eiermann - Universität Karlsruhe, Germany. Iterative methods in Linear Algebra (9/1/92 - 12/31/92).

Paul O. Frederickson - CRAY Research. Parallel algorithms, with emphasis on multigrid algorithms and their application to fluid dynamics and particle models (9/25/91 - 12/31/92).

John Gilbert - Research Scientist, Xerox Palo Alto Research Center. Parallel computing and theoretical computer science (5/1/92 - present).

Elizabeth Jessup - Professor, University of Colorado, Boulder. Parallel matrix computation; eigenvalue problems; programming distributed memory systems (10/19/92 - 12/31/92).

Niel K. Madsen - Lawrence Livermore National Laboratory. Numerical solutions of partial differential equations, with specific interests in method of lines techniques, PDE software, matrix algorithms for vector and parallel computers (10/18/90 - present).

Jeffrey D. McDonald - Applications Specialist, MASPAR Computer Corporation. Computational fluid dynamics (5/1/92 - present).

Alex Pothen - Professor of Mathematics, University of Waterloo, Ontario, Canada. Parallel scalable sparse matrix algorithms (10/01/92 - present).

David Rogers - Project Manager, Molecular Simulations, Inc. Next-generation computational chemistry software product (2/28/92 - 12/31/92).